

VALIDATION OF A CONNECTION BETWEEN ARBITRARY
END-POINTS IN A COMMUNICATIONS NETWORK USING
AN AUGMENTED SPE

CROSS-REFERENCE TO RELATED APPLICATIONS

- 5 This is the first application filed for the present invention.

MICROFICHE APPENDIX

Not Applicable.

TECHNICAL FIELD

- 10 This invention relates to communications networks, and in particular to a method for validating connections established between arbitrary end-points in a communications network.

BACKGROUND OF THE INVENTION

- 15 Within the modern network space, the Synchronous Optical Network (SONET) protocol (and its European equivalent, the Synchronous Digital Hierarchy (SDH)) protocol is a popular mechanism for data transport. Under the SONET/SDH protocol, connections through the network
- 20 core, and between end-user communications devices, are constructed using a layered model. Each layer (or level) uses connections established at lower levels to build connections spanning progressively larger portions of the network. Within the network core, Section and Line-level
- 25 connections are established. Section-level connections (commonly referred to as sections) are set-up between topologically adjacent nodes (which may be, for example, cross-connects, regenerators or Add-Drop-Multiplexors) of

09662337 091400
004160 4229960

the network, and correspond to individual hops of an end-to-end connection. Line-level connections (commonly referred to as lines) are set-up between nodes capable of processing line overhead within data signals (e.g. cross-connects or Add-Drop-Multiplexors). Regenerators are usually capable of terminating only sections, whereas cross-connects and Add-Drop-Multiplexors will normally terminate both sections and lines. Lines typically span one or more sections.

10 An end-to-end connection between nodes at which a data signal is multiplexed and demultiplexed is a path-level connection (usually referred to as a Path). A path commonly spans multiple lines, and may extend beyond the network core to terminate at suitable end-user communications equipment (such as concentrator access switches or service provider routers).

The SONET/SDH protocol provides conventional methods for validating sections, lines and paths. These include section trace, section parity, and line parity checks which implement integrity and validation within one section or one SONET/SDH line. Telecordia TR-253 describes a Path Trace implementation, which enables an end-to-end validation of a path. Path parity implements a path level integrity check. Tandem connection monitoring implements additional integrity checks within the path level. These path-level integrity checks utilize validation, parity and fault indication data that are inserted into Synchronous Payload Envelopes (SPEs) of signals being transported through the path, typically embedded within the path overhead (POH).

Co-pending and co-assigned United States Patent Application No. 09/539,707 filed on March 31, 2000, and

004760" 2E29960

entitled METHOD AND SYSTEM FOR ESTABLISHING
CONTENT-FLEXIBLE CONNECTIONS IN A COMMUNICATIONS NETWORK
teaches a technique for establishing an open connection
(OP-N), mapped across a communications network. The OP-N
5 connection is "concatenatable", in that an end user can
transport arbitrarily concatenated signal traffic through
the OP-N connection. In principle, virtually any
combination of concatenated and non-concatenated signals
may be used, up to the bandwidth capacity of the OP-N
10 connection. The traffic mixture (i.e., the mix of
concatenated and non-concatenated traffic) within the OP-N
connection can be selected by the end user to satisfy their
requirements, and may be changed by the end user as those
requirements change, without requiring re-configuration of
15 the OP-N connection.

The OP-N connection described in the
above-referenced co-pending patent application is
constructed using a layered model, in a manner analogous to
conventional SONET/SDH connections. The layers of an OP-N
20 connection are designed to fit between the SONET/SDH Line
and Path layers. In general, an OP-N connection is
expected to span multiple lines, and yet be shorter than an
end-to-end path. An OP-N connection may carry multiple
paths and will be set up and validated prior to the
25 establishment of any paths through it. It is therefore
necessary to implement a technique for validating OP-N
connections independently of any Path-level connections.

Co-pending and co-assigned United States patent
Application No. 09/597,974, filed on June 20, 2000 and
30 entitled VALIDATION OF A CONNECTION BETWEEN ARBITRARY
END-NODES IN A COMMUNICATIONS NETWORK teaches a technique
for validating OP-N connections independently of any

0966237 091400

Path-level connections. Validation of the OP-N connection is accomplished by inserting performance monitor (PM) information into an unused portion of the transport overhead (TOH) of a SONET/SDH data signal at a transmitting end-node of the OP-N connection. At the receiving end-node of the OP-N connection, the PM information is extracted and examined. At each intermediate node between the transmitting and receiving end-nodes, the PM information is extracted from the data signal, buffered while the data signal is pointer processed, and then reinserted before forwarding the data signal. Multiple levels of OP-N connections can be independently validated in this manner, by providing each level with respective PM information.

A disadvantage of this method is the need for extracting, buffering, and reinserting the PM information at intermediate nodes between the transmitting and receiving end-nodes of the OP-N connection. Such complex processing requires specialized hardware and/or software within each node, which significantly increases costs.

Accordingly, a cost-effective method and apparatus for validating an OP-N connection mapped through a communications network between arbitrary end-points remains highly desirable.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a method and apparatus for ensuring the integrity and validity of a connection mapped through a communications network between arbitrary end-points.

Another object of the present invention is to provide a method and apparatus for ensuring the integrity

and validity of a connection mapped through a communications network between arbitrary end-points, in which the connection is topologically longer than one SONET/SDH line or section, and shorter than a SONET/SDH path.

A further object of the present invention is to provide a method and apparatus for ensuring the integrity and validity of multiple layers of connections that are topologically longer than one SONET/SDH line or section, and shorter than a SONET/SDH path.

Accordingly, an aspect of the present invention provides method of validating a connection mapped through a communications network between first and second end-nodes. At the first end-node, performance monitor (PM) information is inserted into a predetermined location within a payload envelope of a synchronous container. The synchronous container is transmitted through the connection to the second end-node. At the second end-node, the PM information is extracted from the synchronous container.

A further aspect of the present invention provides a system of validating a connection mapped through a communications network between first and second end-nodes. The system comprises: means for inserting performance monitor (PM) information into a predetermined location within a payload envelope of the synchronous container at the first end-node; means for transmitting the synchronous container through the connection to the second end-node; and means for extracting the PM information from the synchronous container at the second end-node.

Another aspect of the present invention provides an apparatus for validating a connection mapped through a

09662337 "091400

communications network between first and second end-nodes. The apparatus comprises: means for inserting performance monitor (PM) information into a predetermined location within a payload envelope of a synchronous container; and
5 means for transmitting the synchronous container through the connection.

Another aspect of the present invention provides an apparatus for validating a connection mapped through a communications network between first and second end-nodes.
10 The apparatus comprising: means for receiving a synchronous container through the connection; and means for extracting performance monitor (PM) information from a predetermined location within a payload envelope of the synchronous container.

15 *949* *A* The synchronous container is preferably an augmented synchronous payload envelope (SPE) having a payload capacity that is sufficient to accommodate the PM information. In such cases, the first and second nodes are preferably enabled to pointer process the augmented SPE.
20 This may be accomplished by providing an extended a range of valid payload pointer values that can be pointer processed within each node participating in the connection.

In some embodiments, each node in the network is adapted to support a plurality of connection layers and the
25 connection is mapped on one of the plurality of connection layers. In such cases, PM information respecting each layer is inserted into a respective predetermined location of the synchronous container.

The PM information may include any one or more of a
30 Trace field; a Parity field; and an indicator field.

004760 252960

The Trace field may include a nibble of a trace message for communicating information concerning the connection.. Successive nibbles of the trace message may be inserted into respective successive synchronous containers
5 until an entire trace message has been sent. The trace message may be repeated after the entire trace message has been sent.

The parity field may contain a parity value calculated in respect of an outgoing synchronous container.
10 The calculated parity value is inserted into the parity field of the next outgoing synchronous container. The parity value may be a BIP-8, and may be calculated starting after an H2 byte of a transport overhead (TOH) portion of the signal, and incorporating all SPE bytes until the H2
15 byte of a next data signal.

The indicator field may contain an error count accumulated in respect of the data signal. The error count may be a BIP-8.

In some embodiments, the step of extracting the PM
20 information comprises a step of extracting one or more of a trace field; a parity field; and an indicator field. Following extraction of the PM information, the synchronous container can be converted into a SONET/SDH SPE having a conventional size and format.

25 A recalculated parity value may be calculated in respect of the received data signal. The recalculated parity value can be compared with a received parity value contained in the extracted parity field to obtain an error count. The error count value in respect of the received
30 data signal may be accumulated.

004637 091400
004637 091400

004637 091400

The step of extracting the indicator field may include: monitoring the indicator field of each successive received data signal; and asserting an AIS state if the indicator field of each of a first predetermined number of successive data signals contains a first predetermined value. The AIS state can be de-asserted if the indicator field of each of the first predetermined number of successive data signals contains a value other than the first predetermined value. The first predetermined number of successive data signals may be three, and the first predetermined value may be binary "1111".

The step of extracting an indicator field may include: monitoring the indicator field of each successive received signal; and asserting an RDI state if the indicator field of each of a second predetermined number of successive signals contains a second predetermined value. The predetermined number of successive data signals may be three, and the second predetermined value may be binary "1100".

20 BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

25 Fig. 1. is a block diagram schematically illustrating a communications network in which an embodiment of the present invention may be implemented;

Fig. 2 is a schematic diagram illustrating an exemplary layered architecture of connections mapped through the communications network of Fig. 1;

30

Fig. 3 is a schematic illustration of an exemplary format of connection performance monitor (PM) overhead for each layer of OP-N connections illustrated in Fig. 2;

Fig. 4 is a schematic illustration of exemplary operations in a method of validating OP-N connections in accordance with an embodiment of the present invention; and

Figs. 5a and 5b are schematic diagrams respectively illustrating exemplary transmit (Tx) and receive (Rx) processors, in accordance with an embodiment of the present invention.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention provides a method and apparatus for ensuring the integrity and validity of multiple layers of open (OP-N) connections mapped through a communications network.

FIG. 1 is a schematic diagram of a communications network 10 in which the present invention may be implemented. The network 10 includes a plurality of nodes which may be, for example, cross-connects 12a-d, regenerators 14a-b or Add-Drop-multiplexors (ADMs) 16a-c. Communications devices 18a, 18b (e.g. concentrator access switches or service provider routers), may be connected to the network 10 via one or more edge nodes 20a, 20b. Each of the cross-connects 12a-d, regenerators 14a-b, ADMs 16a-b, and edge nodes 20a-b are preferably optical nodes configured for wave division multiplex (WDM) and/or dense wave division multiplex (DWDM) transport of data traffic.

As shown in Fig. 1, in order to support arbitrarily concatenated signal traffic between the end-user communications devices 18a and 18b, an OP-N connection 26 is mapped through the network 10 between the edge nodes 20a and 20b. The OP-N connection 26 is set up through "OP-N enabled" nodes, and is adapted to transport arbitrarily concatenated signal traffic up to the maximum bandwidth capacity of the OP-N. The OP-N connection 26 is an OP-48 connection (i.e. $N = 48$, see Fig. 2) so that the OP-N connection 26 has a bandwidth capacity equivalent to an optical carrier OC-48 signal. A SONET/SDH path 28 is mapped between end-user communications devices 18a and 18b through the OP-N connection 26, and utilizes some, or all, of the bandwidth capacity of the OP-N connection 26.

For the purposes of the present invention, an "OP-N enabled" node shall be understood to refer to a network node (of any type) having suitable hardware and software to enable it to participate in an OP-N connection, as described in applicant's above-referenced patent application. An OP-N enabled node is also capable of inserting and processing performance monitor (PM) information in a manner described below in greater detail.

Fig. 2 is a schematic diagram illustrating the construction of the OP-N connection 26 and the end-to-end SONET/SDH path 28. As shown in Fig. 2, both of these connections are constructed using a layered architecture. The SONET/SDH protocol defines section level connections 30a-h which correspond to individual hops in the path 28 between the communications devices 18a and 18b. Line connections 32a-f span one or more sections 30 between network nodes that are capable of processing SONET/SDH line overhead within data signals being transported through the

004760" 2229960

path 28. Both section and line connections 30 and 32 are defined in the SONET/SDH standard. Similarly, the path 28 can be treated as a conventional SONET/SDH path connection because, at least from the point of view of the communications devices 18a and 18b, the path 28 is indistinguishable from a conventional SONET/SDH path. Thus conventional path level verification and integrity checks may be utilized. However, in the illustrated embodiment the path 28 is mapped through the OP-N connection 26. Accordingly, between the edge nodes 20a and 20b (defining respective end-points of the OP-N connection 26), transport of data signals through the path 28 will be handled in accordance with the transport functions of the OP-N connection 26 rather than standard SONET/SDH.

As shown in Fig. 2, the OP-N connection 26 is constructed using a layered model. In the illustrated embodiment, this layered model employs three layers of OP-N connections, in which connections of each successive layer span progressively larger distances (in terms of the number of hops) through the network 10. It will be appreciated that the number of layers of connections is arbitrary, so that OP-N connections can be constructed through a network utilizing a layered model having more or fewer layers than illustrated in Fig. 2.

In the example shown in Fig. 2, high bandwidth layer 1 OP-N connections spanning comparatively short distances are mapped between selected nodes within the network. These layer 1 OP-N connections include a first OP-768 34a mapped between cross-connect 12a and Add-Drop Multiplexor 16a; a second OP-768 connection 34b mapped between Add-Drop Multiplexor 16a and cross-connect 12d; and a pair of feeder OP-192 connections 34c and 34d

respectively linking edge nodes 20a and 20b to cross-connects 12a and 12d.

00460" 4329960
The layer 2 OP-N connections include an OP-48 connection 36a mapped between the edge node 20a and Add-Drop Multiplexor 16a, and utilizing a portion of the bandwidth of the layer 1 OP-768 and OP-192 connections 34a and 34c; and an OP-96 connection 36b mapped between the edge node 20b and the Add-Drop Multiplexor 16a that utilizes a portion of the bandwidth of each of the layer 1 OP-768 and OP-192 connections 34b and 34d. Finally, the OP-N connection 26 is a layer 3 OP-48 connection mapped between the edge nodes 20a and 20b. This connection occupies the entire bandwidth of the layer 2 OP-48 connection 36a, and a portion of the bandwidth of the layer 2 OP-96 connection 36b. In order to ensure reliable data communications using the SONET/SDH path 28 mapped through the OP-N connection 26, the present invention provides a method for validating each of the OP-N connections on each respective layer, independently of one another.

As shown in Fig. 3, validation of OP-N connections in each layer is accomplished by converting a conventional SONET/SDH SPE 40 into an augmented SPE 40'', and inserting performance monitor (PM) information 38 into a predetermined location within the augmented SPE 40''. In principle, the predetermined location within the augmented SPE 40' is arbitrary, and therefore can be selected as required for a specific implementation. However the predetermined location is preferably selected so as to minimize interference or overlap with path overhead and payload data contained in the original (i.e. conventional SONET/SDH) SPE 40. Thus, it is preferable to avoid

0966237 091400

inserting the PM information 38 into the payload overhead (POH) portion 44 of the augmented SPE 40'', because the POH 44 is typically fully defined for the payload data within the payload capacity 46 of the original SPE 40. It is
5 therefore generally not possible to assume that any unused space within the POH 44 will be available for carrying the PM information 38. Similarly, it is preferable to avoid interspersing the PM information within the payload data contained in the payload capacity 46, as this would tend to
10 increase the processing required to insert and extract the PM information 38 into and from the SPE. Accordingly, the PM information 38 is preferably appended onto the payload capacity 46, starting after the last byte of payload data, as shown in Fig. 3. Alternatively, the PM information 38
15 can be prepended to the payload capacity 46, by inserting it between the POH 44 and the payload data.

In order to accommodate the increased size of the augmented SPE 40'', the conventional STS frame 42 is replaced by an expanded frame 42', which contains the
20 conventional transport overhead (TOH) and an envelope capacity that has been enlarged to accommodate the augmented SPE 40''. Proper pointer processing of the expanded frame 42 and augmented SPE 40'' can be accomplished by expanding the range of allowable payload
25 pointer values. In particular, the payload pointer (not shown) within the H2 bytes of each frame 42 is normally provided as a 10-bit binary number, which is capable of accommodating an SPE of up to 1024 bytes. This is significantly larger than the 783 bytes mandated by the
30 SONET/SDH standard. Accordingly, it is possible to define an augmented SPE 40', which is similar to a conventional SPE 40, but has a payload capacity that has been expanded, up to a maximum size of 1024 bytes (per STS-1) to accept

the addition of PM information 38. For example, in the embodiment shown in Fig. 2, three layers of OP-N connections are provided. If the PM information 38 associated with each OP-N layer is 8 bytes in length, then
5 an augmented SPE 40'' having a payload capacity of 807 bytes, for example, will be able accommodate the PM information 38.

Fig. 3 illustrates the conversion of a conventional STS-1 SPE 40 into an augmented SPE 40'' by the addition of
10 a single column, which thereby extends the size of the STS-1 SPE 40 by 9 bytes. It will be appreciated that the present invention is not limited to STS-1 SPE's, but rather may be applied to any required SPE granularity. Thus, for example, an STS-48c SPE 40 could be converted into an
15 augmented STS-48c SPE 40' by the addition of a column, on a per STS-1 basis, to increase the size of the STS-48c SPE 40 by $48 \times 9 = 432$ bytes. Similarly, it will be appreciated that more columns may be added as required, up to a maximum size of 1024 bytes per STS-1.

20 *Sub A27* An augmented SPE 40' can be inserted into a synchronous container (e.g. expanded data frame 42') using conventional means, and transported through an OP-N connection mapped over any transport infrastructure in which the allowable frame length has been increased and
25 pointer processor state machines have been adjusted to accept the increased range of allowable payload pointer values. The conventional 125 μ Sec. frame frequency can be maintained within the OP-N connection by either increasing the line bit-rate, or by adjusting the location of the
30 augmented SPE 40'' within the synchronous container to utilize available unused space within the transport overhead (TOH). For example, in the case of a concatenated

096227 091400

547 A² →

STS-48c signal, the transport overhead of only the first STS-1 frame is required for routing the signal through the end-to-end path 28. Within each of the second through 48th STS-1 frames of the STS-48c signal, a concatenation indicator is inserted into the TOH to indicate that the respective frame is to be treated in an identical manner to that of the previous STS-1 frame. As a result, much of the TOH within each of the 2nd through 48th frames of an STS-48c signal is redundant, and may therefore be utilized for accommodating an augmented SPE 40', at least while the signal is being conveyed through an OP-N connection.

It will be appreciated that the present invention is not limited to SONET/SDH frames, and may be utilized in conjunction with any synchronous container capable of carrying an augmented SONET/SDH SPE 40' across the network fabric.

Conventional SONET/SDH pointer processing techniques operate to pass the SPE 40 of each frame through the node unchanged. This conventional functionality is preferably maintained with the augmented SPE 40'' by enabling pointer processor state machines to accommodate an appropriately increased range of allowable payload pointer values. For example, in the case of three-layers of OP-N connections, and 8 bytes of PM information per OP-N layer, the allowable payload pointer range must be extended from a conventional range of 0-782 to an augmented range of 0-806.

Carrying the PM information within an augmented SPE 40'' permits the PM information 38 to pass transparently through intermediate nodes between the end-nodes of an OP-N connection. The frame TOH can therefore be terminated at intermediate nodes in a conventional manner without

09662337-091400

004760" 2229960

affecting the PM information 38. If desired, the PM information 38 can be read at one or more intermediate nodes to enable auditing and fault isolation functionality. At the receiving end-point of an OP-N connection, the augmented SPE 40' can be restored to the standard SPE format by stripping the PM information 38 from the augmented SPE 40''.

As shown in Fig. 3, PM information, 38a-c in this example, associated with each OP-N connection layer is inserted into a respective predetermined location within the augmented SPE 40''. Thus for a three-layer OP-N architecture, as illustrated in Fig. 2, three blocks of PM information 38a-c are inserted into respective locations within the augmented SPE 40''. Each block of PM information 38a-c is preferably divided into predetermined fields to facilitate various connection integrity and validation checks. Exemplary fields include an OP-N trace 48, OP-N parity 50 and OP-N indicator 52. The use of each of these fields in validation and integrity checks of each OP-N connection is described below in greater detail, with particular reference to Fig. 4.

In use, PM information 38 is inserted into an augmented SPE 40'' by a node serving as a transmitting end-point of an OP-N connection, and extracted from the augmented SPE 40''' and examined by a node serving as a receiving end-point of the OP-N connection (see Fig. 4). At nodes intermediate the end-points of an OP-N connection, the augmented SPE 40'' (including the PM information 38) is pointer processed through the node without change. If desired, the PM information can be read at nodes intermediate the end-points of an OP-N connection to enable auditing and fault isolation.

For example, in the case of a signal being transported through the path 28 between the communications devices 18a and 18b, a conventional SPE is received at the end node 20a from the communications device 18a. Within the end-node 20a, the conventional SPE is converted into an augmented SPE 40'' (e.g. by increasing its size by 24 bytes) to accept PM information 38 respecting three layers of OP-N connections. Layer 1, 2 and 3 PM information 38a, b, c (Fig. 3) in respect of the OP-192 connection 34c, OP-48 connection 36a and OP-48 connection 26, respectively, is then prepared and inserted into the augmented SPE 40'' by the edge node 20a. At the cross-connect 12a, the layer 1 PM information 38a is extracted from the signal and examined, while the layer 2 and layer 3 PM information 38b and 38c is pointer processed through the cross-connect 12a within the augmented SPE 40''. At the regenerator 14a, layer 1, layer 2 and layer 3 PM information 38a-c is pointer-processed through the node with the augmented SPE 40''. A new layer 1 PM information 38a in respect of the OP-768 connection 34a is then inserted into the augmented SPE 40'' by the cross-connect 12a prior to transmitting the signal through the OP-768 connection 34a to the Add-Drop Multiplexor 16a. At Add-Drop Multiplexor 16a, layer 1 and layer 2 PM information 38a,b is extracted and examined, while the layer 3 PM information 38c is pointer processed through the node with the augmented SPE 40'. The Add-Drop Multiplexor 16a then prepares and inserts new layer 1 and layer 2 PM information 38a,b respecting the OP-768 connection 34b and the OP-96 36b, prior to transmitting the frame. The OP-768 connection 34b includes one intermediate node (regenerator 14b). The OP-96 connection 36b includes two intermediate nodes (regenerator 14b and cross-connect 12d). At the intermediate node 14b the

004760" 2EE29960

layer 1, layer 2 and layer 3 PM information 38a-c is pointer processed through the node along with the augmented SPE 40'. When the signal arrives at the cross-connect 12d, the layer 1 PM information 38a respecting the OP-768 connection 34b (Fig. 2) is extracted from the signal and examined while the layer 2 and layer 3 PM information 38b and 38c is pointer processed through the cross-connect 12d with the augmented SPE 40'. New layer 1 PM information 38a respecting the feeder OP-192 connection 34d is prepared within the cross-connect 12d and inserted into the signal, prior to launching the signal toward the edge node 20b, which extracts and examines the PM information 38a, 38b and 38c to validate each of layer 1 OP-192 connection 34d, layer 2 OP-96 connection 36b, and layer 3 OP-48 connection 26. Following extraction of the PM information 38a-c by the end-node 20b, the augmented SPE 40' is converted back into a conventional SONET/SDH SPE 40 for transmission to the receiving communications device 18b.

As described above, at a node serving as a transmitting end-node of an OP-N connection (e.g. edge node 20a in the example of Fig. 2), PM information 38 (including, for example, OP-N trace, parity and indicator fields) respecting the OP-N connection are prepared and inserted into an augmented SPE 40'. At a node serving as a receiving end-node of the OP-N connection (e.g. end node 20b in the example of Fig. 2), the respective PM information 38 is extracted from the augmented SPE 40' and examined to determine the validity and integrity of the OP-N connection. At any pointer processing nodes intermediate the transmitting and receiving end-nodes of the OP-N connection, the PM information 38 is within the augmented SPE 40' without change.

007650' 2229960

The purpose and function of each of the OP-N trace, parity, and indicator fields 48-52 of the PM information 38 is explained below with reference to Fig. 4. As will be understood by those skilled in the art, an OP-N connection
5 is normally enabled in both directions between the end nodes. However, for the sake of clarity, transmission in only one direction is illustrated.

OP-N Trace

The OP-N trace field 48 of the PM information 38
10 carries a trace message (having a length of, for example, 16 bytes) carrying connectivity information for a respective OP-N connection. The trace message is preferably divided into smaller parts (e.g. nibbles having a length of, for example, 4 bits), with successive nibbles
15 being included in the PM information 38 inserted into successive augmented SPEs of a data signal. If a 16-byte trace message is divided into 4 bit nibbles, the trace message repeats every 32 SPE's. The format of the trace message can be arbitrary. It is normally defined by a
20 software application that is used to implement the OP-N PM information messaging functionality.

At the transmitting end-node of an OP-N connection, a trace message in respect of the OP-N connection is prepared for insertion into outgoing frames. At the
25 receiving end-node of the OP-N connection, the trace message is extracted from received augmented SPEs 40' and examined to verify the connectivity of the OP-N connection. If the trace message is sent in 4 bit nibbles, each nibble is stored in memory in the receiving node until the
30 entire 16 byte message has been accumulated and reassembled.

OP-N Parity

00460" 2EE2960

The OP-N parity field 52 of the PM information 38 provides a parity check between transmit and receive nodes in an OP-N connection. For example, the hardware within each node can process and report OP-N parity on a per-frame or per-signal (e.g. STS-48) basis. If desired, OP-N parity may be used to generate Errored and Severely Errored seconds counts rather than being used to infer precise bit error ratios.

OP-N parity can advantageously be calculated as an 8-bit, Bit Interleaved Parity (BIP-8) value using even parity (so that the parity will be correct during AIS insertion, described below), over a predetermined portion of a data signal. For example, a BIP-8 value can be calculated starting after the H2 byte of TOH covering all STS-N SPE bytes until the next H2 byte. Positive and negative stuff (H3) bytes, as well as the corresponding OP-N trace and OP-N indicator fields 48 and 50, may also be included. The BIP-8 result calculated for an outgoing frame can be inserted into the OP-N parity field 52 of the next frame.

At a transmitting end-node of an OP-N connection, OP-N parity is calculated for outgoing frames (and inserted into the parity field 52 of the augmented SPE 40' contained within the next outgoing frame). At the receiving end-node of the OP-N connection, the OP-N parity information is extracted and simultaneously recalculated as successive frames are received. These extracted and recalculated values can then be compared to assess the integrity of the OP-N connection.

OP-N Indicator Field

The indicator field 50 of the PM information 38 enables fault information to be conveyed to the end-nodes

of the OP-N connection. Various fault indications may be provided in respect of each OP-N connection. Exemplary fault indications include Remote Error Indication (REI), Alarm Indication Signaling (AIS), and Remote Defect Indication (RDI). Each of these are described in greater detail below.

OP-N Remote Error Indication (REI)

OP-N Remote Error Indication (REI) provides a capability to signal bit errors to the far-end on a frame-by-frame basis.

In the receive direction, each node can accumulate a valid OP-N error count (e.g. a BIP-8 error count based on received and recalculated parity values, as described above) as successive frames are received. At a node serving as a transmitting end-node of an OP-N connection, the node hardware can insert the OP-N error count accumulated as the frames were received. At the receiving end-node of the OP-N connection, the error count can be accumulated as successive frames are received.

OP-N Alarm Indication Signaling (AIS)

OP-N Alarm Indication Signaling (AIS) is used to signal a failure downstream to the receiving end-node of the OP-N connection. When either Line or Path AIS is asserted (in accordance with conventional SONET/SDH protocol) at any node participating in an OP-N connection, the payload as well as the Transparent Container is set to all 1's.

At each end-node in the OP-N connection, OP-N AIS is asserted if the OP-N REI error count information contains "1111" for 3 consecutively received frames. OP-N AIS is de-asserted if a value, other than "1111", remains

004T60" 2529960

unchanged for 3 consecutive frames. When OP-N AIS is asserted (or when otherwise enabled by software), a value of "1100" is inserted into the OP-N indicator field 48 of the PM information of each outgoing frame.

5 OP-N Remote Defect Indication (RDI)

OP-N Remote Defect Indication (RDI) is used to signal conditions such as OP-N-AIS and Trace mismatch to the far-end of the OP-N. OP-N RDI uses the same set of bits that is used for OP-N REI information.

10 At each end-node in the OP-N connection, OP-N RDI is asserted if the OP-N REI information contains "1100" for 3 consecutively received frames. OP-N RDI is de-asserted if a value, other than "1100", is received unchanged for 3 consecutive frames.

15 Fig. 5a is a block diagram illustrating an exemplary implementation of an OP-N Transmit (Tx) processor 54 for processing data signals including PM information 38 in accordance with the present invention. The OP-N Tx processor 54 may be implemented (by hardware and software) within each OP-N enabled node of the network 10 capable of serving as a transmitting end-node of an OP-N connection (e.g. edge node 20a in the example of Fig. 2). In general, the Tx processor 54 comprises an elastic store 56; a Tx framer 58; a memory 60; and a pointer-processor state machine 62. The elastic store 56, which may be implemented as a first-in-first-out (FIFO) having an adjustable read pointer (not shown) operates under the control of the pointer processor state machine 62 to map an incoming data stream (which may, for example, be conventional SONET/SDH SPE's) into augmented SPEs 40 containing the OP-N PM information 38. The Tx framer 58 operates to insert the augmented SPEs 40 into outgoing

004T60" 25229960

frames of a data signal for transmission through an OP-N connection.

The PM information 38 can be inserted into successive augmented SPEs 40' on the read side of the elastic store 56. In order to insert the PM information 38 into an augmented SPE 40', the applicable values of the OP-N trace field 48 and OP-N indicator field 50 are read from the memory 60, and inserted into the augmented SPE 40'. For the OP-N parity field 52, an outgoing parity value is calculated in respect of a previous outgoing frame, and the result is inserted into the augmented SPE 40'. This operation results in the calculation of a parity value for each outgoing frame, with the calculated value being inserted in the augmented SPE 40' of the next outgoing frame.

Fig. 5b is a block diagram illustrating an exemplary implementation of a OP-N receive (Rx) processor 64 for processing incoming frames including PM information 38 in accordance with the present invention. The OP-N Rx processor 64 may be implemented (with hardware and software) within each OP-N enabled node of the network 10, and generally comprises an Rx framer 66; an elastic store 68; a memory 70; and a pointer-processor state machine 72. The Rx framer 66 operates to receive successive frames of an incoming data signal, and extracts the frame payload (e.g. the respective augmented SPEs 40'). The augmented SPEs 40' (including the PM information 38) are forwarded to the elastic store 68, which may be implemented as a first-in-first-out (FIFO) having an adjustable read pointer (not shown). Processing of the augmented SPE 40' through the elastic store 68 is controlled by the pointer processor state machine 72. Thus

004660" 2E29960

the pointer processor state machine 72 controls the process of reading data from the elastic store 68 such that the PM information 38 is extracted from the augmented SPE 40', which is then restored to the standard SONET/SDH SPE size and format.

09652337 091400

The PM information 38 extracted from augmented SPEs 40' is processed as described above with reference to Fig. 4. Thus the contents of the OP-N trace field 48 is saved in the memory 70. The contents of the OP-N indicator field 50 is processed by a fault indication processor 74 (i.e. for OP-N REI, OP-N AIS and OP-N RDI) and saved in the memory 70. An accumulated BIP-8 OP-N REI error count can also be saved in the memory 70. The contents of the OP-N parity field 52 is supplied to an input of a parity comparator 76 as a received parity value. A recalculated parity value, which is determined by a parity calculation circuit 78 on the basis of received frames of the incoming data signal, is supplied to another input of the parity comparator 76. The parity comparator 76 operates to compare the received and recalculated parity values, and saves the comparison result in the memory 70 as an accumulated parity error value.

The "receiver-side" processing of received frames described above is implemented in a receiving end-node of an OP-N connection (e.g. edge node 20b in the illustration of Fig. 2). If desired, this processing can also be implemented in intermediate nodes for auditing and fault isolation purposes. In either case, "receiver-side" processing results in the accumulation of parity and REI error values in the memory 70 to permit evaluation of the integrity and validity of the OP-N connection.

The present invention therefore provides a method and apparatus for validating OP-N connections mapped through a communications network.

The embodiment(s) of the invention described above
5 is(are) intended to be exemplary only. The scope of the invention is therefore intended to be limited solely by the scope of the appended claims.

004760" 2E29960